

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims

1. (Canceled)

2. (Currently amended) The method of claim [[1]] 6, wherein the storage system has at least two processing nodes, wherein the plurality of logical processing units are distributed across the at least two processing nodes, wherein one processing node includes a plurality of central processing units, and wherein in the event of the failure of the first logical processing unit, the plurality of processing nodes stay operational.

3. (Currently amended) The method of claim [[1]] 6, wherein an administrative console is coupled to the plurality of processing nodes of the storage system, further comprising: prior to partitioning, processing, at the administrative console, information on processing requirements, memory requirements and host bus adapter requirements for the plurality of logical processing units.

4. (Currently amended) The method of claim [[1]] 6, wherein one or more partitioning applications are coupled to the plurality of logical processing units, further comprising:

in response to grouping the at least two logical processing units, starting initial program load of the first logical processing unit;

determining via the one or more partitioning applications an identification of the second logical processing unit grouped with the first logical processed unit; and

presenting, by the one or more partitioning applications, common resources to the first and second logical processing units.

5. (Currently amended) The method of claim [[1]] 6, further comprising: receiving from the first logical processing unit, a request for memory access of a logical processing unit;

determining, by the one or more partitioning applications coupled to the plurality of logical processing units, whether the logical processing unit is grouped with the first logical processing unit;

if the logical processing unit is grouped with the first logical processing unit, then allowing the memory access of the logical processing unit to the first logical processing unit; and

if the logical processing unit is not grouped with the first logical processing unit, then preventing the memory access of the logical processing unit to the first logical processing unit.

6. (Currently amended) ~~The method of claim 1, further comprising:~~ A method, comprising:

partitioning a plurality of processing nodes in a storage system into a plurality of logical processing units, wherein the plurality of logical processing units can respond to I/O requests from a host coupled to the storage system;

grouping at least two logical processing units, wherein data in a first storage coupled to a first logical processing unit of the least two logical processing units is mirrored by data in a second storage coupled to the second logical processing unit of the at least two logical processing units; and

in response to a failure of the first logical processing unit, responding to an I/O request from the host via the second logical processing unit, wherein the method further comprises:

(i) receiving from the first logical processing unit, a request for a reinitialized program load of a logical processing unit;

(ii) determining, by one or more partitioning applications coupled to the plurality of logical processing units, whether the logical processing unit is grouped with the first logical processing unit;

(iii) if the logical processing unit is grouped with the first logical processing unit, then allowing the reinitialized program load of the logical processing unit; and

(iv) if the logical processing unit is not grouped with the first logical processing unit, then preventing the reinitialized program load of the logical processing unit.

7. (Currently amended) The method of claim [[1]] 6, further comprising:

receiving a write request from the host to the plurality of processing nodes in the storage system; and

writing, by the one or more partitioning applications, data corresponding to the write request to the first storage coupled to the first logical processing unit and the second storage coupled to the second logical processing unit.

8. (Currently amended) The method of claim [[1]] 10, further comprising:
receiving a read request from the host to the plurality of processing nodes in the storage system; and

reading, by one or more partitioning applications, data corresponding to the read request from the first storage coupled to the first logical processing unit.

9. (Currently amended) The method of claim [[1]] 10, wherein the partitioning and grouping are performed by one or more partitioning applications coupled to the plurality of processing nodes, wherein the one or more partitioning applications comprise a hypervisor application of a redundant system.

10. (Currently amended) ~~The method of claim 1;~~ A method, comprising:
partitioning a plurality of processing nodes in a storage system into a plurality of logical processing units, wherein the plurality of logical processing units can respond to I/O requests from a host coupled to the storage system;
grouping at least two logical processing units, wherein data in a first storage coupled to a first logical processing unit of the least two logical processing units is mirrored by data in a second storage coupled to the second logical processing unit of the at least two logical processing units; and
in response to a failure of the first logical processing unit, responding to an I/O request from the host via the second logical processing unit, wherein the partitioning and grouping further comprises:

(i) associating first pool numbers sequentially to a first set of logical processing units included in a first processing node;

(ii) associating second pool numbers sequentially to a second set of logical processing units included in a second processing node; and

(iii) assigning pairs of logical processing units with the same associated pool numbers to be partner virtual machines, wherein the partner virtual machines mirror each other and provide redundancy.

11. (Canceled)

12. (Currently amended) The system of claim [[11]] 16, wherein the storage system has at least two processing nodes, wherein the plurality of logical processing units are distributed across the at least two processing nodes, wherein one processing node includes a plurality of central processing units, and wherein in the event of the failure of the first logical processing unit, the plurality of processing nodes stay operational.

13. (Currently amended) The system of claim [[11]] 16, further comprising:
an administrative console coupled to the plurality of processing nodes of the storage system; and
means for processing, at the administrative console, information on processing requirements, memory requirements and host bus adapter requirements for the plurality of logical processing units.

14. (Currently amended) The system of claim [[11]] 16, further comprising:
one or more partitioning applications coupled to the plurality of logical processing units;
means for starting initial program load of the first logical processing unit, in response to grouping the at least two logical processing units;
means for determining via the one or more partitioning applications an identification of the second logical processing unit grouped with the first logical processed unit; and
means for presenting, by the one or more partitioning applications, common resources to the first and second logical processing units.

15. (Currently amended) The system of claim [[11]] 16, further comprising:

means for receiving from the first logical processing unit, a request for memory access of a logical processing unit;

means for determining, by the one or more partitioning applications coupled to the plurality of logical processing units, whether the logical processing unit is grouped with the first logical processing unit;

means for allowing the memory access of the logical processing unit to the first logical processing unit if the logical processing unit is grouped with the first logical processing unit; and

means for preventing the memory access of the logical processing unit to the first logical processing unit if the logical processing unit is not grouped with the first logical processing unit.

16. (Currently amended) ~~The system of claim 11, further comprising:~~ A system, comprising:

a storage system;

a plurality of processing nodes in the storage system;

means for partitioning the plurality of processing nodes in the storage system into a plurality of logical processing units, wherein the plurality of logical processing units can respond to I/O requests from a host coupled to the storage system;

means for grouping at least two logical processing units, wherein data in a first storage coupled to a first logical processing unit of the least two logical processing units is mirrored by data in a second storage coupled to the second logical processing unit of the at least two logical processing units; and

means for responding to an I/O request from the host via the second logical processing unit in response to a failure of the first logical processing unit, wherein the system further comprises:

(i) means for receiving from the first logical processing unit, a request for a reinitialized program load of a logical processing unit;

(ii) means for determining, by one or more partitioning applications coupled to the plurality of logical processing units, whether the logical processing unit is grouped with the first logical processing unit;

(iii) means for allowing the reinitialized program load of the logical processing unit; if the logical processing unit is grouped with the first logical processing unit; and

(iv) means for preventing the reinitialized program load of the logical processing unit if the logical processing unit is not grouped with the first logical processing unit.

17. (Currently amended) The system of claim [[11]] 16, further comprising:
means for receiving a write request from the host to the plurality of processing nodes in the storage system; and
means for writing, by the one or more partitioning applications, data corresponding to the write request to the first storage coupled to the first logical processing unit and the second storage coupled to the second logical processing unit.

18. (Currently amended) The system of claim [[11]] 20, further comprising:
means for receiving a read request from the host to the plurality of processing nodes in the storage system; and
means for reading, by one or more partitioning applications, data corresponding to the read request from the first storage coupled to the first logical processing unit.

19. (Currently amended) The system of claim [[11]] 20, wherein the partitioning and grouping are performed by one or more partitioning applications coupled to the plurality of processing nodes, wherein the one or more partitioning applications comprise a hypervisor application of a redundant system.

20. (Currently amended) ~~The system of claim 11, A system, comprising:~~
a storage system;
a plurality of processing nodes in the storage system;
means for partitioning the plurality of processing nodes in the storage system into a plurality of logical processing units, wherein the plurality of logical processing units can respond to I/O requests from a host coupled to the storage system;
means for grouping at least two logical processing units, wherein data in a first storage coupled to a first logical processing unit of the least two logical processing units is mirrored by data in a second storage coupled to the second logical processing unit of the at least two logical processing units; and

means for responding to an I/O request from the host via the second logical processing unit in response to a failure of the first logical processing unit, wherein the partitioning and grouping further comprises:

(i) means for associating first pool numbers sequentially to a first set of logical processing units included in a first processing node;

(ii) means for associating second pool numbers sequentially to a second set of logical processing units included in a second processing node; and

(iii) means for assigning pairs of logical processing units with the same associated pool numbers to be partner virtual machines, wherein the partner virtual machines mirror each other and provide redundancy.

21. (Canceled)

22. (Currently amended) The ~~article of manufacture~~ computer readable storage medium of claim [[21]] 26, wherein the storage system has at least two processing nodes, wherein the plurality of logical processing units are distributed across the at least two processing nodes, wherein one processing node includes a plurality of central processing units, and wherein in the event of the failure of the first logical processing unit, the plurality of processing nodes stay operational.

23. (Currently amended) The ~~article of manufacture~~ computer readable storage medium of claim [[21]] 26, wherein an administrative console is coupled to the plurality of processing nodes of the storage system, the operations further comprising:

prior to partitioning, processing, at the administrative console, information on processing requirements, memory requirements and host bus adapter requirements for the plurality of logical processing units.

24. (Currently amended) The ~~article of manufacture~~ computer readable storage medium of claim [[21]] 26, wherein one or more partitioning applications are coupled to the plurality of logical processing units, the operations further comprising:

in response to grouping the at least two logical processing units, starting initial program load of the first logical processing unit;

determining via the one or more partitioning applications an identification of the second logical processing unit grouped with the first logical processed unit; and

presenting, by the one or more partitioning applications, common resources to the first and second logical processing units.

25. (Currently amended) ~~The article of manufacture~~ computer readable storage medium of claim [[21]] 26, the operations further comprising:

receiving from the first logical processing unit, a request for memory access of a logical processing unit;

determining, by the one or more partitioning applications coupled to the plurality of logical processing units, whether the logical processing unit is grouped with the first logical processing unit;

if the logical processing unit is grouped with the first logical processing unit, then allowing the memory access of the logical processing unit to the first logical processing unit; and

if the logical processing unit is not grouped with the first logical processing unit, then preventing the memory access of the logical processing unit to the first logical processing unit.

26. (Currently amended) ~~The article of manufacture of claim 21, A computer~~ readable storage medium including code, wherein the code when executed by a processor causes operations, the operations comprising:

partitioning a plurality of processing nodes in a storage system into a plurality of logical processing units, wherein the plurality of logical processing units can respond to I/O requests from a host coupled to the storage system;

grouping at least two logical processing units, wherein data in a first storage coupled to a first logical processing unit of the least two logical processing units is mirrored by data in a second storage coupled to the second logical processing unit of the at least two logical processing units; and

in response to a failure of the first logical processing unit, responding to an I/O request from the host via the second logical processing unit, the operations further comprising:

(i) receiving from the first logical processing unit, a request for a reinitialized program load of a logical processing unit;

(ii) determining, by one or more partitioning applications coupled to the plurality of logical processing units, whether the logical processing unit is grouped with the first logical processing unit;

(iii) if the logical processing unit is grouped with the first logical processing unit, then allowing the reinitialized program load of the logical processing unit; and

(iv) if the logical processing unit is not grouped with the first logical processing unit, then preventing the reinitialized program load of the logical processing unit.

27. (Currently amended) The ~~article of manufacture~~ computer readable storage medium of claim [[21]] 26, the operations further comprising:

receiving a write request from the host to the plurality of processing nodes in the storage system; and

writing, by the one or more partitioning applications, data corresponding to the write request to the first storage coupled to the first logical processing unit and the second storage coupled to the second logical processing unit.

28. (Currently amended) The ~~article of manufacture~~ computer readable storage medium of claim [[21]] 30, the operations further comprising:

receiving a read request from the host to the plurality of processing nodes in the storage system; and

reading, by one or more partitioning applications, data corresponding to the read request from the first storage coupled to the first logical processing unit.

29. (Currently amended) The ~~article of manufacture~~ computer readable storage medium of claim [[21]] 30, wherein the partitioning and grouping are performed by one or more partitioning applications coupled to the plurality of processing nodes, wherein the one or more partitioning applications comprise a hypervisor application of a redundant system.

30. (Currently amended) ~~The article of manufacture of claim 21~~ A computer readable storage medium including code, wherein the code when executed by a processor causes operations, the operations comprising:

partitioning a plurality of processing nodes in a storage system into a plurality of logical processing units, wherein the plurality of logical processing units can respond to I/O requests from a host coupled to the storage system;

grouping at least two logical processing units, wherein data in a first storage coupled to a first logical processing unit of the least two logical processing units is mirrored by data in a second storage coupled to the second logical processing unit of the at least two logical processing units; and

in response to a failure of the first logical processing unit, responding to an I/O request from the host via the second logical processing unit, wherein the partitioning and grouping further comprises:

(i) associating first pool numbers sequentially to a first set of logical processing units included in a first processing node;

(ii) associating second pool numbers sequentially to a second set of logical processing units included in a second processing node; and

(iii) assigning pairs of logical processing units with the same associated pool numbers to be partner virtual machines, wherein the partner virtual machines mirror each other and provide redundancy.